Synergistic Computation of Planar Maxima and Convex Hull

Jérémy Barbay[1] and Carlos Ochoa[1⋆]

Departamento de Ciencias de la Computación, Universidad de Chile, Chile
jeremy@barbay.cl, cochoa@dcc.uchile.cl.

arXiv:1702.08545v1 [cs.DS] 27 Feb 2017

Abstract. Refinements of the worst case complexity over instances of fixed input size consider the input order or the input structure, but rarely both at the same time. Barbay et al. [2016] described "synergistic" solutions on multisets, which take advantage of the input order and the input structure, such as to asymptotically outperform any comparable solution which takes advantage only of one of those features. We consider the extension of their results to the computation of the Maxima Set and the Convex Hull of a set of planar points. After revisiting and improving previous approaches taking advantage only of the input order or of the input structure, we describe synergistic solutions taking optimally advantage of various notions of the input order and input structure in the plane. As intermediate results, we describe and analyze the first adaptive algorithms for Merging Maxima and Merging Convex Hulls.

Keywords: Convex Hull, Dominance Query, Maxima, Membership Query, Multivariate Analysis, Synergistic.

1 Introduction

One way to close the gap between practical performance and the worst case complexity over instances of fixed input size is to refine the later, considering smaller classes of instances. Such measures of difficulty can be seen along two axis: some depend of the *structure* of the input, such as the repetitions in a multiset [19], or the positions of the input points in the plane [1] or in higher dimensions [13]; while some others depend on the *order* in which the input is given, such as for permutations [18] but also for points in the plane [2, 16].

Barbay et al. [5] described various "synergistic" solutions on multisets, which take advantage of both the input structure and the input order, in such a way that each of their solutions is never asymptotically slower than other solutions taking advantage of a subset of these features, but also so that on some family of instances, each of their solution performs an order of magnitude faster than any solution taking advantage of only a subset of those features. They left open the generalization of their results to higher dimensions.

In the context of the computation of the Maxima Set and of the Convex Hull, various refinements of the worst case complexity over instances of fixed size have been known for some time. Kirkpatrick and Seidel described algorithms optimal in the worst case over instances of fixed input and output size, first in 1985 for the computation of the Maxima Set of points in any dimension [13] and then in 1986 for the computation of the Convex Hull in the plane [14]: such results can be classified as focused on the **input structure**, and were further refined in 2009 by Afshani et al. [1] in 2 and 3 dimensions. Following a distinct approach, Levcopoulos et al. [16] in 2002, and Ahn and Okamoto [2] in 2011, studied the computation of the Convex Hull

⋆ Corresponding author.

in conjunction with various notions of **input order**: these results can be generalized to take advantage of the input order when computing Maxima Sets (Section 3.1) and can be further refined using recent techniques (Section 4.1). Yet no algorithm (beyond a trivial dovetailing combination of the solutions described above) is known to take advantage of both the **input structure** and **input order** at the same time for the computation of the Maxima Set or of the Convex Hull of points, in the plane or in higher dimension, nor for any other problem than Sorting Multisets.

Hypothesis. It seems reasonable to expect that Barbay et al.'s synergistic results [5] on Sorting Multisets should generalize to similar problems in higher dimension, such as the computation of the Maxima Set and of the Convex Hull of a set of points in the plane. Yet these two problems present new difficulties of their own: (1) while the results on multisets [5] are strongly based on a variant of Demaine et al.'s instance optimal algorithm to Merge Multisets [10], at this date no such results are known for Merging Maxima Sets, and the closest known result for Merging Convex Hulls [3] (from 2008) is not adaptive to the size of the output, and hence not adaptive to the structure of the instance; furthermore (2) whereas many *input order adaptive* results are known for Sorting Multisets (with two surveys in 1992 on the topic [11, 18], and various additional results [4, 24] since then), it seems that none are known for the computation of the Maxima Set and that only a few are known for the computation of the Convex Hull [2, 16].

Our Results. After reviewing previous results on the computation of Convex Hull taking advantage of either the input structure (Section 2.1) or the input order (Section 2.2), and one result on Sorting Multisets taking advantage of both (Section 2.3), we confirm the hypothesis by (1) presenting new solutions for Merging Maxima Sets (Section 3.2) and Merging Convex Hulls (Section 4.2) in the plane, (2) defining new techniques to take advantage of the input order to compute Maxima Sets (Section 3.1), (3) improving previous techniques to analyze the computation of Convex Hulls in function of the input order (Section 4.1), and (4) synthesizing all those results in synergistic algorithms to compute Maxima Sets (Section 3.3) and Convex Hulls (Section 4.3) of a set of planar points. For the sake of pedagogy, **we present our results incrementally, from the simplest to the most complex**. We define formally the notions of input order, input structure and synergistic solution in Section 2. We then describe synergistic solutions for the computation of both the Maxima Set (Section 3) and of the Convex Hull (Section 4) of points in the plane (the latter requiring more advance techniques). In both case, our solution is based on an algorithm merging several partial solutions (Sections 3.2 and 4.2), adapted from Barbay et al.'s `Quick Hull Merge` algorithm [5] to merge multisets[1]. We conclude in Section 5 with a partial list of issues left open for improvement. Due to space constraints we only state our results in the article and defer all the proofs to the appendix.

2 Background

Beyond the worst case complexity over instances of fixed size, the adaptive analysis of algorithms refines the scope of the analysis by considering the worst case complexity over finer classes of instances. We describe here some relevant results along two axis: results about the computation of the Convex Hull which take advantage of the *input structure* (Section 2.1), results about the computation of the Convex Hull which take advantage of the *input order* (Section 2.2), and one result about Sorting Multisets which depends of both (Section 2.3).

[1] Itself inspired from Demaine et al.'s algorithm solving the same problem [9].

2.1 Input Structure

In 1985, Kirkpatrick and Seidel [13] described an algorithm to compute the Maxima Set of points in any dimension, which is optimal in the worst case over instances of input size n and output size h, running in time within $O(n \log h)$. One year later, in 1986, they [14] described a slightly more complex algorithm to compute the Convex Hull in the plane, which is similarly optimal in the worst case over instances of input size n and output size h, running in time within $O(n \log h)$. Both results are described as *output sensitive*, in the sense that the complexity depends on the size of the output, and can be classified as adaptive to the *input structure*, as the position of the points clearly determine the output (and its size), as opposed to algorithms described in the next paragraph taking advantage of the order in which those points are given.

2.2 Input Order

A *polygonal chain* is a curve specified by a sequence of points $p_1, \ldots, p_n$. The curve itself consists of the line segments connecting the pairs of consecutive points. A polygonal chain C is *simple* if any two edges of C that are not adjacent are disjoint, or if the intersection point is a vertex of C; and any two adjacent edges share only their common vertex. Melkman [17] described an algorithm that computes the Convex Hull of a simple polygonal chain in linear time, and Chazelle [8] described an algorithm for testing whether a polygonal chain is simple in linear time. In 2002, Levcopoulos et al. [16] combined these results to yield an algorithm for computing the Convex Hull of polygonal chains. Their algorithm tests if the chain C is simple, using Chazelle's algorithm [8]: if the chain C is simple, the algorithm computes the Convex Hull of C in linear time, using Melkman's algorithm [17]. Otherwise, if C is not simple, the algorithm partitions C into the subsequences C' and C'', whose sizes differ at most in one; recurses on each of them; and merges the resulting Convex Hulls using Preparata and Shamos's algorithm [22]. They measured the complexity of this algorithm in terms of the minimum number of simple subchains κ into which the chain C can be partitioned. Let $t(n, \kappa)$ be the worst-case time complexity taken by this algorithm for an input chain of n vertices that can be partitioned into κ simple subchains. They showed that $t(n, \kappa)$ satisfies the following recursion relation: $t(n, \kappa) \leq t(\lceil \frac{n}{2} \rceil, \kappa_1) + t(\lfloor \frac{n}{2} \rfloor, \kappa_2), \kappa_1 + \kappa_2 \leq \kappa + 1$. The solution to this recursion gives $t(n, \kappa) \in O(n(1 + \log \kappa)) \subseteq O(n \log n)$.

In 2011, Ahn and Okamoto [2] followed a distinct approach for the computation of the Convex Hull, also based on some notions of input order. They considered a variant of the problem where the output is the same size of the input, but such that the Convex Hull can be checked and extracted in linear time from this output. In this context, they describe adaptive results directly inspired from disorder measures introduced for the study of adaptive algorithms for Sorting Permutations, such as `Runs` and `Inv` [11, 18].

Inspired by Ahn and Okamoto's definition [2], we define some simple measure of *input order* for the computation of Maxima Sets in Section 3.1, and we slightly refine Levcopoulos et al.'s analysis [16] for the computation of Convex Hulls in Section 4.1.

2.3 Synergistic Solutions

Inspired by previous results on sorting multisets in a way adaptive to the frequencies of the element [19] on one hand, and on sorting permutation in a way adaptive to the distribution of the lengths of the subsequences of consecutive positions already sorted [24] on the other hand, Barbay et al. [5] described two "synergistic" algorithms Sorting Multisets, which take advantage both of the input structure and of the input order, in such a way that each of their solutions is never asymptotically slower than other solutions taking advantage of a subset of these

features, but also so that on some family of instances, each of their solution performs an order of magnitude faster than any solution taking advantage of only a subset of those features. They left open the generalization of their results to higher dimensions. We generalize their results to dimension 2, for the computation of the Maxima Set in Section 3, and for the computation of the Convex Hull in Section 4.

3 Maxima Set

Given a point $p \in \mathbb{R}^2$, let p_x and p_y denote the x- and y-coordinates of p, respectively. Given two points p and q, p *dominates* q if $p_x \geq q_x$ and $p_y \geq q_y$. Given a set $\mathcal{S}$ of points in d dimensions, a point p from $\mathcal{S}$ is called *maximal* if none of the other points of $\mathcal{S}$ dominates p. The Maxima Set of such a set $\mathcal{S}$ is the uniquely defined set of all maximal points [15]. Kirkpatrick and Seidel [13] described an algorithm that computes the Maxima Set running in time within $O(n \log h)$, where n is the number of input points, and h is the number of points in the Maxima set (i.e., the size of the output). In 2009, Afshani et al. [1] improved the results on the computation of both the Maxima Set and Convex Hull in dimension 2 and 3 by taking the best advantage of the relative positions of the points (while ignoring the input order).

If the points are sorted by their coordinates (say, in lexicographic order of their coordinates for a fixed order of the dimensions), the Maxima Set can be computed in time linear in the size of the input. Refining this insight, we show in Section 3.1 that one can take advantage of the input order even if it is not as strictly sorted: this presents a result which is orthogonal to previous input structure adaptive results [1, 13]. In order to combine these results synergistically with previous input structure adaptive results [1, 13], we study in Section 3.2 an algorithm that solves the problem of Merging Maxima, which asks for computing the Maxima Set of the union of maxima sequences, in such a way that it outperforms both input structure adaptive results by taking advantage of the number and sizes of the maxima sequences and of the relative positions between the points in them. Last, we combine those results into a single synergistic algorithm, which decomposes the input sequence of points into several "easy" subsequences for which the corresponding Maxima Set can be computed in linear time, and then proceeds to merge them. The resulting algorithm not only outperforms previous input structure adaptive results [1, 13] in the sense that it never performs (asymptotically) worse, and performs better when it can take advantage of the input order, it also outperforms a dovetailing combination of previous input structure adaptive algorithms [1, 13] and the input order adaptive algorithm described in Section 3.1.

3.1 Input Order Adaptive Maxima Set

In many cases the Maxima Set can be computed in time linear in the size of the input, independently from the size of the Maxima Set itself. For instance, consider an order of the input where (1) the maximal points are given in order sorted by one coordinates, and (2) for each maximal point p, all the points dominated by p are given immediately after p in the input order (in any relative order). The Maxima Set of a sequence of points given in this order can be extracted and validated in linear time by a simple greedy algorithm, which throws an exception if the input is not in such an order. Each of the various ways to deal with such exceptions directly yields an input order adaptive algorithm [2]. For instance, if the point found to be out of order is *inserted* in the partial Maxima Set computed up to this point, this yields an algorithm running in time within $O(n \lg \texttt{Inv})$ where `Inv` is the sum of such insertion costs.

Let's label such a sequence "*smooth*", and by extension any input subsequence of consecutive positions which have the same property. Given an input sequence $\mathcal{S}$, let σ denote the minimal

number of smooth subsequences into which it can be decomposed. Most interestingly for synergistic purpose, such a decomposition can be computed in time linear in the input size. Detecting such σ smooth subsequences and merging them two by two yields an algorithm running in time within $O(n(1 + \log \sigma))$. Such a result is orthogonal to previous input structure adaptive results [1, 13]: it can be worse than $O(n \log h)$ when the output size h is small and the input is in a "bad" order, as it can be much better than $O(n \log h)$ when h is large and the input is in a "good" order. We show in the next two sections an algorithm which outperforms both.

3.2 Union of Maxima Sequences

We describe the `Quick Union Maxima` algorithm, which computes the Maxima Set of the union of maxima sequences in the plane, assuming that the points in the maxima sequences are given in sorted order by their x-coordinates (i.e., Merging Maxima). This algorithm generalizes Barbay et al.'s [5] `QuickSort` inspired algorithm for Merging Multiset and is a building block towards the synergistic algorithm for computing the Maxima Set of a set of planar points described in Section 3.3. Given a maxima sequence $\mathcal{M}_i$, let $\mathcal{M}_i[a]$ and $\mathcal{M}_i[b..c]$ denote the a-th point and the block of consecutive $c - b + 1$ points corresponding to positions from b to c in $\mathcal{M}_i$, respectively. As its name indicates, the algorithm is inspired by the `QuickSort` algorithm.

Description of the algorithm Quick Union Maxima. The `Quick Union Maxima` algorithm chooses a point p that forms part of the Maxima Set of the union, and discards all the points dominated by p. Note that all dominated points do not belong to the maxima that contains p. The selection of p ensures that at least half of the maxima sequences will have points dominated by p or to the right of p and at least half of the maxima sequences will have points dominated by p or to the left of p. The algorithm identifies a block $\mathcal{B}$ of consecutive points in the maxima sequence that contains p, which forms part of the output Maxima Set (p is contained in $\mathcal{B}$). All the points in $\mathcal{B}$ are discarded. Note that if the points of a maxima sequence in the plane are sorted in ascending order by their x-coordinates, then their y-coordinates are sorted in decreasing order. This algorithm takes advantage of this fact when it discards points. All discarded points are identified by doubling searches [6] inside the maxima sequences . The algorithm then recurses separately on the non-discarded points to the left of p and on the non-discarded points to the right of p. (See Algorithm 1 for a more formal description.) Next, we analyze the time complexity of the `Quick Union Maxima` algorithm.

Complexity Analysis of Quick Union Maxima. Every algorithm for Merging Maxima needs to certify that blocks of consecutive points in the maxima sequences are dominated or are in the Maxima Set of the union. In the following we formalize the notion of a *certificate*, that permits to check the correctness of the output in less time than to recompute the output itself. We define a "language" of basic "arguments" for such certificates: *domination* (which discards points from the input) and *maximality* (which justify the presence of points in the output) arguments, and their key positions in the instance. A certificate will be verified by checking each of its arguments: those can be checked in time proportional to the number of blocks in them.

Definition 1. $\langle \mathcal{M}_i[a] \supset \mathcal{M}_j[b..c] \rangle$ *is an* elementary domination argument *if the point* $\mathcal{M}_i[a]$ *dominates all the points in the block* $\mathcal{M}_j[b..c]$.

Definition 2. $\langle \mathcal{M}_i[a] \supset \mathcal{M}_{j_1}[b_1..c_1], \ldots, \mathcal{M}_{j_t}[b_t..c_t] \rangle$ *is a* Domination Argument *if the point* $\mathcal{M}_i[a]$ *dominates all the points in the blocks* $\mathcal{M}_{j_1}[b_1..c_1], \ldots, \mathcal{M}_{j_t}[b_t..c_t]$.

Algorithm 1 `Quick Union Maxima`

Input: A set $\mathcal{M}_1, \ldots, \mathcal{M}_\rho$ of ρ maxima sequences
Output: The Maxima Set of the union of $\mathcal{M}_1, \ldots, \mathcal{M}_\rho$

1: Compute the median μ of the x-coordinates of the middle points of the maxima sequences;
2: Perform doubling searches for the value μ in the x-coordinates of the points of all maxima sequences, starting at both ends of the maxima sequences in parallel;
3: Find the point p of maximum y-coordinate among the points q such as $q_x \geq \mu$ in all maxima sequences, note $j \in [1..\rho]$ the index of the maxima sequence containing p;
4: Discard all points dominated by p through doubling searches for the values of p_x and p_y in the x- and y-coordinates of all maxima, respectively, except $\mathcal{M}_j$ (Search for p_x in the points q such that $q_x \geq \mu$ and for p_y in the points q such that $q_x < \mu$);
5: Find the point r of maximum y-coordinate among the points q such that $q_x > p_x$ in all maxima sequences except $\mathcal{M}_j$ and find the point ℓ of maximum x-coordinate among the points q such that $q_y > p_y$ in all maxima sequences except $\mathcal{M}_j$;
6: Discard the block in $\mathcal{M}_j$ containing p that forms part of the output through doubling searches for the values of ℓ_x and r_y in the x- and y-coordinates of the points in $\mathcal{M}_j$, respectively. (Search for ℓ_x in the points q such that $q_x < p_x$ and for r_y in the points q such that $q_x > p_x$.);
7: Recurse separately on the non-discarded points left and right of p.

Lemma 1. *A* domination argument $\langle \mathcal{M}_i[a] \supset \mathcal{M}_{j_1}[b_1..c_1], \ldots, \mathcal{M}_{j_t}[b_t..c_t] \rangle$ *can be checked in* $O(t)$ *data comparisons.*

It is not enough to eliminate all points that can not participate in the output. Certifying would still require additional work: a correct algorithm must justify the optimality of its output. To this end we define maximality arguments.

Definition 3. $\langle \mathcal{M}_i[a..b] \dashv \mathcal{M}_{j_1}[a_1..b_1], \ldots, \mathcal{M}_{j_t}[a_t..b_t] \rangle$ *is a* Maximality Argument *if* ***either*** *the points* $\mathcal{M}_i[b]$ *dominates the points* $\mathcal{M}_{j_1}[a_1], \ldots, \mathcal{M}_{j_t}[a_t]$ *and the* x*-coordinates of the points* $\mathcal{M}_{j_1}[a_1 - 1], \ldots, \mathcal{M}_{j_t}[a_t - 1]$ *are less than the* x*-coordinate of the point* $\mathcal{M}_i[a]$ ***or*** *the point* $\mathcal{M}_i[a]$ *dominates the points* $\mathcal{M}_{j_1}[b_1], \ldots, \mathcal{M}_{j_t}[b_t]$ *and the* y*-coordinates of the points* $\mathcal{M}_{j_1}[b_1 + 1], \ldots, \mathcal{M}_{j_t}[b_t + 1]$ *are less than the* y*-coordinate of the point* $\mathcal{M}_i[b]$*.*

If $\langle \mathcal{M}_i[a..b] \dashv \mathcal{M}_{j_1}[a_1..b_1], \ldots, \mathcal{M}_{j_t}[a_t..b_t] \rangle$ is a valid *maximality argument*, then the points in the block $\mathcal{M}_i[a..b]$ are maximal among the maxima sequences $\mathcal{M}_i, \mathcal{M}_{j_1}, \ldots, \mathcal{M}_{j_t}$.

Lemma 2. *A* maximality argument $\langle \mathcal{M}_i[a..b] \dashv \mathcal{M}_{k_1}[a_1..b_1], \ldots, \mathcal{M}_{k_t}[a_t..b_t] \rangle$ *can be checked in* $O(t)$ *data comparisons.*

The difficulty of finding and describing domination and maximality arguments depend on the points they refer to in the maxima sequences, a notion captured by "argument points":

Definition 4. *Given an argument* $\mathcal{A} = \langle \mathcal{M}_i[a] \supset \mathcal{M}_{j_1}[b_1..c_1], \ldots, \mathcal{M}_{j_t}[b_t..c_t] \rangle$ *or* $\mathcal{B} = \langle \mathcal{M}_i[a..b] \dashv \mathcal{M}_{j_1}[a_1..b_1], \ldots, \mathcal{M}_{j_t}[a_t..b_t] \rangle$*, the* Argument Points *are the points* $\mathcal{M}_i[a]$ *in* $\mathcal{A}$ *and* $\mathcal{M}_i[a]$ *and* $\mathcal{M}_i[b]$ *in* $\mathcal{B}$*.*

Those atomic arguments combine into a general definition of a certificate that any correct algorithm for Merging Maxima in the comparison model can be modified to output.

Definition 5. *Given a set of maxima sequences and their* Maxima Set $\mathcal{M}$ *expressed as several blocks on the maxima sequences. A* certificate *of* $\mathcal{M}$ *is a set of domination and maximality arguments such that the* Maxima Set *of any instance satisfying those arguments is given by the description of* $\mathcal{M}$*. The length of a certificate is the number of distinct argument points in it.*

We divide the analysis of the time complexity of the `Quick Union Maxima` algorithm into two lemmas. We first bound the cumulated time complexity of the doubling searches for the value of the median μ of the x-coordinates of the middle points of the maxima (i.e., step 2 of Algorithm 1) and the doubling searches in the points discard steps (i.e., steps 4 and 6 of Algorithm 1). The algorithm partitions the maxima sequences into blocks of consecutive discarded points, where each block is discarded because it is dominated, or because this block forms part of the MAXIMA SET of the union. Each block forms part of some argument of the certificate computed by the algorithm.

Lemma 3. *Let $s_1, \ldots, s_\beta$ be the sizes of the β blocks into which the algorithm `Quick Union Maxima` divides the maxima sequence $\mathcal{M}_i$. The cumulated time complexity of the doubling searches for the value of the medians μ of the x-coordinates of the middle points of the maxima sequences (i.e., step 2) and the doubling searches in the points discard steps (i.e., steps 4 and 6) of the algorithm `Quick Union Maxima` in the maxima sequence $\mathcal{M}_i$ is within $O(\sum_{j=1}^{\beta} \log s_j)$.*

Proof. Every time the algorithm finds the insertion rank of one of the medians μ of the x-coordinates of the middle points of the maxima sequences in $\mathcal{M}_i$, it finds a position d inside a blocks whose points will be discarded. The discard points steps that search for the insertion rank of p_x and p_y start the search in the position d. The time complexity of both discard points steps are bounded by $O(\log s_b)$, where s_b is the size of the discarded block b. Both discard points steps partition $\mathcal{M}_i$ in positions separating the blocks to the left of b, the block b itself and the blocks to the right of b.

The combination of the doubling search that finds the insertion rank of μ in the x-coordinates of $\mathcal{M}_i$ with the doubling searches that discard points starting in d can be represented as a tree. Each internal node has two children, which correspond to the two subproblems into which the recursive steps partition $\mathcal{M}_i$, the blocks to the left of b and the blocks to the right of b . The cost of this combination is bounded by $O(\log s_b + \log s)$, where s is the minimum between the sum of the sizes of the blocks to the left of b and the sum of the sizes of the blocks to the right of b, because of the two doubling searches in parallel. The size of each internal node is the size of the block discarded in this step. The size of each leaf is the sum of the sizes of the blocks in the child subproblem represented by this leaf.

We prove that at each combination of steps, the total cost is bounded by eight times of the sum of the logarithms of the sizes of the nodes in the tree. This is done by induction over the number of steps. If the number of steps is zero then there is no cost. For the inductive step, if the number of steps increases by one, then a new combination of steps is done and a leaf subproblem is partitioned into two new subproblems. At this step, a leaf of the tree is transformed into an internal node and two new leaves are created. Let w and z such that $w \leq z$ be the sizes of the new leaves created. Note that w and z are the sum of the sizes of the blocks to the left and to the right, respectively, of the discarded block b in this step. The cost of this step is less than $4 \log w + 4 \log b$. The cost of all the steps then increases by $4 \log w + 4 \log b$, and hence eight times the sum of the logarithms of the nodes in the tree increases by $8(\log w + \log z + \log b - \log(w + z + b))$. But if $w \geq 3$, $w \geq b$ and $w \leq z$ then the result follows. □

We bound next the time complexity of the steps that compute the median μ of the x-coordinates of the middles points in the maxima (i.e., step 1 of Algorithm 1) and the steps that find the points p, ℓ and r (i.e., steps 3 and 5 of Algorithm 1) in the `Quick Union Maxima` algorithm. Note that one execution of these steps has time complexity bounded by the number of maxima sequences in the sub-instance. The partition of the maxima sequences by the x-coordinate of p and the discarded points decrease the number of maxima sequences in the subinstances.

Lemma 4. *The cumulated number of comparison performed by the steps that compute the median μ of the x-coordinates of the middles points in the ρ maxima sequences (i.e., step 1) and the steps that find the points p, ℓ and r (i.e., steps 3 and 5) in the* `Quick Union Maxima` *algorithm is within $O(\sum_{i=1}^{\delta} \log \binom{\rho}{m_i})$, where δ is the length of the certificate $\mathcal{C}$ computed by the algorithm and $m_1, \ldots, m_\delta$ is a sequence where m_i is the number of maxima sequences whose blocks form the i-th argument of $\mathcal{C}$.*

Proof. We prove this lemma by induction over δ and ρ. The time complexity of one of these steps is linear in the number of maxima sequences in the sub-instance (i.e., ignoring all the empty maxima sequences of this sub-instance).

Let $\mathcal{T}(\delta, k)$ be the cumulated time complexity during the execution of the steps that compute the medians μ of the x-coordinates of the middles points (i.e., step 1) and during the steps that find the points p, ℓ and r (i.e., steps 3 and 5) in the algorithm `Quick Union Maxima`. We prove that $\mathcal{T}(\delta, k) \leq \sum_{i=1}^{\delta} m_i \log \frac{k}{m_i} - k$, where m_i is the number of maxima sequences whose blocks form the i-th argument of $\mathcal{C}$. Let μ be the first median of the x-coordinates of the middles points of the maxima sequences computed by the algorithm. Let c and d be the number of maxima sequences that have non-discarded point only above of p_y and to the right of p_x, respectively. Let b be the number of maxima sequences that have non-discarded points above p_y and to the right of p_x. Let e be the number of maxima sequences which all their points are dominated by p. Let δ_c and δ_d be the number of arguments computed by the algorithm to discard points in the maxima sequences above p_y and to the right of p_x, respectively. Then, $\mathcal{T}(\delta, k) = \mathcal{T}(\delta_c, c+b) + \mathcal{T}(\delta_d, d+b) + k$ because of the two recursive calls and the steps 1, 3 and 5 of the algorithm `Quick Union Maxima`. By Induction Hypothesis, $\mathcal{T}(\delta_c, c+b) \leq \sum_{i=1}^{\delta_c} m_i \log \frac{c+b}{m_i} - c - b$ and $\mathcal{T}(\delta_d, d+b) \leq \sum_{i=1}^{\delta_d} m_i \log \frac{d+b}{m_i} - d - b$. In the worst case $e = 0$, and we need to prove that $c + d \leq \sum_{i=1}^{\delta_c} m_i \log \left(1 + \frac{d}{c+b}\right) + \sum_{i=1}^{\delta_d} m_i \left(1 + \frac{c}{d+b}\right)$, but this is a consequence of $\sum_{i=1}^{\delta_c} m_i \geq c + b, \sum_{i=1}^{\delta_d} m_i \geq d + b$ (the number of discarded blocks is greater than or equal to the number of maxima sequences); $c \leq d + b, d \leq c + b$ (at least $\frac{k}{2}$ maxima sequences are left to the left and to the right of μ); and $\log \left(1 + \frac{y}{x}\right)^x \geq y$ for $y \leq x$. □

Combining Lemma 3 and Lemma 4 yields an upper bound on the number of data comparisons performed by the algorithm `Quick Union Maxima`:

Theorem 1. *Given ρ maxima sequences. The* `Quick Union Maxima` *algorithm performs within $O(\sum_{j=1}^{\beta} \log s_j + \sum_{i=1}^{\delta} \log \binom{\rho}{m_i})$ data comparisons when it computes the* Maxima Set *of the union of these maxima sequences; where β is the number of blocks in the certificate $\mathcal{C}$ computed by the algorithm; $s_1, \ldots, s_\beta$ are the sizes of these blocks; δ is the length of $\mathcal{C}$; and $m_1, \ldots, m_\delta$ is a sequence where m_i is the number of maxima sequences whose blocks form the i-th argument of $\mathcal{C}$. This number of comparisons is optimal in the worst case over instances formed by ρ maxima sequences that have certificates $\mathcal{C}$ formed by β blocks of sizes $s_1, \ldots, s_\beta$ and length δ such that $m_1, \ldots, m_\delta$ is a sequence where m_i is the number of maxima sequences whose blocks form the i-th argument of $\mathcal{C}$.*

The optimality of this algorithm is a consequence of the fact that it checks each arguments of any certificate using a constant number of argument points. We prove next that the `Quick Union Maxima` algorithm computes a *certificate* which length is a constant factor of the length of the certificate of minimal length. Consider the following algorithm for Merging Maxima. The `Left-to-Right` algorithm chooses the first points from left to right of each maxima sequence and computes the points u and v of maximum and second maximum y-coordinate among these points, respectively. Let $\mathcal{M}_i$ be the maxima sequence that contains u. Let a be the index of u in $\mathcal{M}_i$. The y-coordinates of the points in the input maxima are sorted in decreasing order from left to

right. The algorithm searches then for the insertion rank of v_y in $\mathcal{M}_i$. Let b be the index of the rightmost point g in $\mathcal{M}_i$ such that $g_y > v_y$. The block $\mathcal{M}_i[a..b]$ form part of the MAXIMA SET of the union and are discarded by the `Left-to-Right` algorithm. If g dominates v, the algorithm discards all points in the input dominated by g. The algorithm restarts the computation on the non-discarded points.

Lemma 5. *The `Left-to-Right` algorithm computes a certificate of minimal length when it computes the union of ρ maxima sequences.*

Proof. Let u and v be the points with maximum and second maximum y-coordinate among the first points from left to right of the non-discarded points of the maxima sequences. Let $\mathcal{M}_i$ be the maxima sequence that contains u. Then, all points in $\mathcal{M}_i$ with y-coordinate greater than v_y could be discarded (i.e., form part of the output) and v is the point in the maxima sequences that allows to discard the greatest number of consecutive points including u in $\mathcal{M}_i$. Let g and h be consecutive points in $\mathcal{M}_i$ such that $g_y > v_y > h_y$. If g dominates v, then g is the rightmost point in $\mathcal{M}_i$ that dominates v. Hence, g is the point in $\mathcal{M}_i$ that dominates the maximum number of consecutive points including v in the maxima sequence that contains v. These two arguments are enough to prove that the algorithm computes the minimum number of *argument points*. □

Let a be the number of *argument points* of the certificate computed by the `Left-to-Right` algorithm. The number of *argument points* of the certificate that the algorithm `Quick Union Maxima` computes is within a constant factor of a.

Lemma 6. *The algorithm `Quick Union Maxima` computes a certificate which length is a constant factor of the length of the certificate of minimal length.*

Proof. Suppose that there is a block b of consecutive points in a maxima sequence that the `Left-to-Right` algorithm discards because it identifies that the points in b are in the output. Suppose that the algorithm `Quick Union Maxima` running in the same input computes a point p (p is the point of maximum y-coordinate among the points of x-coordinate greater than μ, then p is in the output) that is contained in b. Let r be the point of maximum y-coordinate among the points with x-coordinate greater than p_x computed by the `Quick Union Algorithm`. Let h be the *argument point* used by the `Left-to-Right` algorithm to identify the rightmost point in b. Hence, $r_y < h_y$. Let ℓ be the point of maximum x-coordinate among the points with y-coordinate greater than p_y computed by the `Quick Union Maxima` algorithm. Let u be the *argument point* used by the `Left-to-Right` algorithm to discard dominated points before the identification of b. Hence, ℓ is the same point as u. So, the algorithm `Quick Union Maxima` discards at least the block b using a constant number of *argument points*. The result follows. □

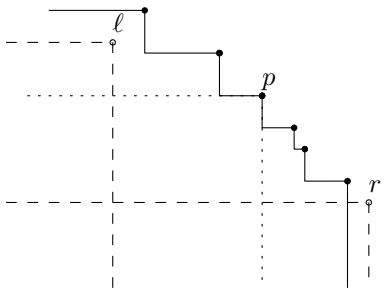

Fig. 1: A representation of a state of the `Quick Union Algorithm` where the points p, ℓ and r has been computed.

In the following section we describe a synergistic result that combines the results of Sections 3.1 and 3.2. This result introduces the synergistic technique also used in the computation of the CONVEX HULL in Section 4.3.

3.3 Synergistic Computation of Maxima Sets

The `(Smooth,Structure) Synergistic Maxima` algorithm decomposes the input of planar points into the minimal number σ of smooth subsequences of consecutive positions, computes their maxima sequences and then merges them using the `Quick Union Maxima` algorithm described in the previous section.

Theorem 2. *Let $\mathcal{S}$ be a set of points in the plane such that $\mathcal{S}$ can be divided into σ smooth maxima sequences. Let h be the number of points in the* Maxima Set *of $\mathcal{S}$. There exists an algorithm that performs within $2n + O(\sum_{j=1}^{\beta} \log s_j + \sum_{i=1}^{\delta} \log \binom{\sigma}{m_i}) \subseteq O(n \log(\min(\sigma, h)))$ [2] data comparisons when it computes the* Maxima Set *of $\mathcal{S}$; where β and $s_1, \ldots, s_\beta$ are the number and sizes of the blocks in the certificate $\mathcal{C}$ computed by the union algorithm, respectively; δ is the length of $\mathcal{C}$; and $m_1, \ldots, m_\delta$ is a sequence where m_i is the number of maxima sequences whose blocks form the i-th argument of $\mathcal{C}$. This number of comparisons is optimal in the worst case over instances $\mathcal{S}$ formed by σ smooth sequences which* Maxima Set *have certificates $\mathcal{C}$ of length δ formed by β blocks of sizes $s_1, \ldots, s_\beta$, such that $m_1, \ldots, m_\delta$ is a sequence where m_i is the number of maxima sequences whose blocks form the i-th argument of $\mathcal{C}$.*

We prove that the number of comparisons performed by the algorithm `(Smooth, Structure) Synergistic Maxima` is asymptotically optimal in the worst case over instances formed by n points grouped in σ smooth sequences, with a final Maxima Set of size h. The upper bound is a consequence of the Theorem 1 and the linear time partitioning algorithm described in Section 3.1. We describe the intuition for the lower bound below: it is a simple adversary argument, based on the definition of a family of "hard" instances for each possible value of the parameters of the analysis, building over each other.

First, we verify the lower bound for "easy" instances, of finite difficulty: general instances formed by a single ($\sigma = 1$) smooth sequence obviously require $\Omega(n)$ comparisons (no correct algorithm can afford to ignore a single point of the input, which could dominate all others), while general instances dominated by a single point ($h = 1$) also require $\Omega(n)$ comparisons (similarly to the computation of the maximum of an unsorted sequence). Each of this lower bound yields a distribution of instances, either of smoothness $\sigma = 1$ or of output size $h = 1$, such that any deterministic algorithm performs $\Omega(n)$ comparisons on average on a uniform distribution of those instances.

Such distributions of "elementary" instances can be duplicated so that to produce various distributions of elementary instances; and combined so that to define a distribution of harder instances:

Lemma 7. *Given the positive integers $n, \sigma, \beta, s_1, \ldots, s_\beta, \delta, m_1, \ldots, m_\delta$, there is a family of instances which can each be partitioned into σ smooth sequences such that,*

- *β and $s_1, \ldots, s_\beta$ are the number and sizes of the blocks in the certificate $\mathcal{C}$ computed by the union algorithm, respectively;*
- *δ is the length of $\mathcal{C}$;*
- *$m_1, \ldots, m_\delta$ is a sequence where m_i is the number of maxima sequences of the smooth sequences whose blocks form the i-th argument of $\mathcal{C}$; and*
- *on average on a uniform distribution of these instances, any algorithm computing the* Maxima Set *of S in the comparison model performs within $\Omega(n + \sum_{i=1}^{\delta} \log \binom{\sigma}{m_i})$ comparisons.*

Finally, any such distribution with a computational lower bound on average yields a computational lower bound for the worst case instance complexity of any randomized algorithm, on

[2] The quantity $\sum_{j=1}^{\beta} \log s_j$ is within $O(n)$ but is much smaller for "easy" instances.

average on its randomness; and as a particular case a lower bound on the worst case complexity of any deterministic algorithm:

Corollary 1. *Given the positive integers $\sigma, \beta, s_1, \ldots, s_\beta, \delta, m_1, \ldots, m_\delta$, and an algorithm A computing the* Maxima Set *of a sequence of n planar points in the comparison model (whether deterministic or randomized), there is an instance I such that A performs a number of comparisons within $\Omega(n + \sum_{i=1}^{\delta} \log \binom{\sigma}{m_i})$ when it computes the* Maxima Set *of I.*

Proof. A direct application of Yao's minimax principle [20, 23, 25]. □

The histories of the computation of the Maxima Set and of the computation of the Convex Hull are strongly correlated: most of the results on one problem also generalize on the other one. Our results on the computation of the Maxima Set similarly generalize to the computation of the Convex Hull, albeit it requires additional work and concepts, which we describe in the next section.

4 Convex Hull

Given a set $\mathcal{S}$ of planar points, the Convex Hull of $\mathcal{S}$ is the smallest convex set containing $\mathcal{S}$ [22]. Given n points in the plane, the problem of computing their Convex Hull is well studied: the worst case complexity over instances of size n is within $\Theta(n \lg n)$ in the algebraic decision tree computational model [21]. Several refinements of this analysis are known: some taking advantage of the input structure [1, 13] and some taking advantage of the input order [2, 16].

Levcopoulos et al. [16] described how to *partition* a sequence of points into subsequences of consecutive positions for which the Convex Hull can be computed in linear time. We refine their analysis to take into account the distribution of the sizes of the subsequences (Section 4.1). This notion of input order for the computation of the Convex Hull is less restrictive than the one seen for the computation of the Maxima Set, in the sense that it allows to consider more sophisticated sequences as "easy" sequences.

As the computation of Convex Hulls reduces to the computation of Upper Hulls (the computation of the Lower Hull is symmetric and completes it into the computation of the Convex Hull), we focus on the latter. We describe an algorithm Merging Upper Hulls in Section 4.2, which yields a synergistic algorithm taking advantage of both the input order and the input structure in Section 4.3. This synergistic algorithm outperforms both the algorithms described by Levcopoulos et al. [16] and Afshani et al. [1], as well as any dovetailing combination of them.

4.1 Input Order Adaptive Convex Hull

A *polygonal chain* is a curve specified by a sequence of points $p_1, \ldots, p_n$. The curve itself consists of the line segments connecting the pairs of consecutive points. A polygonal chain is *simple* if any two edges of P that are not adjacent are disjoint, or if the intersection point is a vertex of P; and any two adjacent edges share only their common vertex. Levcopoulos et al. [16] described an algorithm that computes the Convex Hull of n planar points in time within $O(n \log \kappa)$, where κ is the minimal number of simple chains into which the input sequence of points can be partitioned. The algorithm partitions the points into simple subchains, computes their Convex Hulls, and merges them. In their analysis the complexity of both the partitioning and merging steps are within $O(n \log \kappa)$. In Section 4.3, we describe a partitioning algorithm running in linear time, which is key to the synergistic result.

For a given polygonal chain, there can be several partitions into simple subchains of minimum size for it. We describe a refined analysis which takes into account the relative imbalance between the sizes of the subchains. The idea behind the refinement is to bound the number of operations that the algorithm executes for every simple subchain. This analysis makes it possible to identify families of instances where the complexity of the algorithm is linear even though the number of simple subchains into which the chain is split is logarithmic. In the recursion tree of the execution of the algorithm described by Levcopoulos *et al.* [16] on an input C formed by n planar points, every node represents a subchain of C. The cost of every node is linear in the size of the subchain that it represents. The simplicity test and the merge process are both linear in the number of points in the subchain. Every time this algorithm discovers that the polygonal chain is simple, the corresponding node in the recursion tree becomes a leaf.

Theorem 3. *Given a sequence S of n planar points which can be partitioned into κ simple subchains of respective sizes $n_1, \ldots, n_\kappa$, Levcopoulos et al.'s algorithm [16] computes the convex hull of S in time within $O(n(1 + \mathcal{H}(n_1, \ldots, n_\kappa))) \subseteq O(n(1+\log \kappa)) \subseteq O(n \log n)$, where $\mathcal{H}(n_1, \ldots, n_\kappa) = \sum_{i=1}^{\kappa} \frac{n_i}{n} \log \frac{n}{n_i}$, which is worst-case optimal over instances of n points that can be partitioned into κ simple subchains of sizes $n_1, \ldots, n_\kappa$.*

Proof. Fix the subchain c_i of size n_i. In the worst case, the algorithm considers the n_i points of c_i for the simplicity test and the merging process, in all the levels of the recursion tree from the first level to the level $\lceil \log \frac{n}{n_i} \rceil + 1$, because the sizes of the subchains in these levels are greater than n_i. In the next level, one of the nodes ℓ of the recursion tree fits completely inside c_i and therefore it becomes a leaf. Hence, at least $\frac{n_i}{4}$ points from c_i are discarded for the following iterations. The remaining points of c_i are in the left or the right ends of subchains represented by nodes in the same level of ℓ in the recursion tree. In all of the following levels, the number of operations of the algorithm involving points from c_i can be bounded by the size of the subchains in those levels. So, the sum of the number of the operations in these levels is within $O(n_i)$. As a result, the number of operations of the algorithm involving points from c_i is within $O(n_i \log \frac{n}{n_i} + n_i)$. In total, the time complexity of the algorithm is within $O(n + \sum_{i=1}^{\kappa} n_i \log \frac{n}{n_i}) = O(n(1 + \mathcal{H}(n_1, \ldots, n_\kappa))) \subseteq O(n(1+\log \kappa)) \subseteq O(n \log n)$.

We prove the optimality of this complexity in the worst-case over instances of n points that can be partitioned into κ simple subchains of sizes $n_1, \ldots, n_\kappa$ by giving a tight lower bound. Barbay and Navarro [4] showed a lower bound of $\Omega(n(1 + \mathcal{H}(r_1, \ldots, r_\rho)))$ in the comparison model for Sorting a sequence of n numbers, in the worst case over instances covered by ρ runs (increasing or decreasing) of sizes $r_1, \ldots, r_\rho$, respectively, summing to n. The Sorting problem can be reduced in linear time to the problem of computing the Convex Hulls of a chain of n planar points that can be partitioned into ρ simple subchains of sizes $r_1, \ldots, r_\rho$, respectively. For each real number r, this is done by producing a point with (x, y)-coordinates (r, r^2). The ρ runs (alternating increasing and decreasing) are transformed into ρ simple subchains of the same sizes. The sorted sequence of the numbers can be obtained from the Convex Hull of the points in linear time. □

Similarly to the computation of the Maxima Set, we define in the following section an algorithm for Merging Upper Hulls. This algorithm is a building block towards the synergistic algorithm that computes the Convex Hull of a set of planar points, and is more complex than that for Merging Maxima.

4.2 Union of Upper Hulls

We describe the `Quick Union Hull` algorithms which computes the Upper Hull of the union of ρ upper hull sequences in the plane assuming that the upper hull sequences are given in sorted

Algorithm 2 `Quick Union Hull`

Input: A set $\mathcal{U}_1, \ldots, \mathcal{U}_\rho$ of ρ upper hull sequences
Output: The Upper Hull of the union of the set $\mathcal{U}_1, \ldots, \mathcal{U}_\rho$

1: Compute the median μ of the slopes of the middle edges of the upper hull sequences;
2: Find the point p that has a supporting line of slope μ through doubling searches for the value μ in the slopes of the edges of all upper hull sequences, starting at both ends in parallel, note $j \in [1..\rho]$ the index of the upper hull sequence containing p;
3: Perform doubling searches for the value p_x in the x-coordinates of the points of all upper hull sequences except $\mathcal{U}_j$, starting at both ends in parallel;
4: Find the two tangents of p with all upper hull sequences: the one to the left of p and the one to the right of p, through doubling searches testing for each point q the slope of the two edge that have q as end point and the slope of the line pq and discard the points below these tangents.
5: Discard a block in $\mathcal{U}_j$ containing p that form part of the output by computing the tangent between $\mathcal{U}_j$ and the upper hull sequences left of p of minimum slope and the tangent between $\mathcal{U}_j$ and the upper hull sequences right of p of maximum slope.
6: Repeat until there is no more than one upper hull sequence of size 1: pair those left of p and pair those right of p and apply the Step 4 to those pairs;
7: Discard all points that lie below the lines that joins p with the leftmost point and the rightmost point of the upper hull sequences;
8: Recurse on the non-discarded points left and right of p.

order by their x-coordinates. Given an upper hull sequence $\mathcal{U}_i$, let $\mathcal{U}_i[a]$ and $\mathcal{U}_i[b..c]$ denote the a-th point and the block of $c - b + 1$ consecutive points corresponding to the positions from a to b in $\mathcal{U}_i$, respectively. Given two points p and q, let $m(p, q)$ denote the slope of the straight line that passes trough p and q.

Description of the algorithm Quick Union Hull. The `Quick Union Hull` algorithm is inspired by an algorithm described by Chan et al. [7]. It chooses an edge of slope μ from the upper hull sequences, and computes the point p that has a supporting line of slope μ. The algorithm then splits the points in the upper hull sequences by p_x. It computes the two tangents of p with all the upper hull sequences: the one to the left of p and the one to the right of p, and discards all the points below these tangents. The algorithm also computes a block of consecutive points in the upper hull sequence that contains p which points are part of the output, and discards the points in this block (p is in this block). (This last step is key to the optimality of the algorithm and is significantly more complex than its counterpart in the Merging Maxima solution.) The algorithm then recurses on the non-discarded points to the left of p and on the non-discarded points to the right of p. All these steps take advantage of the fact that the points in the upper hull sequences are sorted in increasing order of their x-coordinates and the slopes of the edges of the upper hull sequences are monotonically decreasing from left to right. (See Algorithm 2 for a more formal description.)

In the following we describe in more details the Step 5 of Algorithm 2. We describe only how to compute a block of consecutive points to the right of p in $\mathcal{U}_j$ that form part of the output, as the left counterpart is symmetric. Let τ be the tangent of maximum slope between $\mathcal{U}_j$ and the upper hull sequences to the right of p. Let q be the point in $\mathcal{U}_j$ that lies in τ. Let λ be the tangent of maximum slope among those computed at Step 4. All the points to the right of p are below λ and λ is a separating line between the portion of $\mathcal{U}_j$ that contains q and the points to the right of p. Given two upper hull sequences $\mathcal{U}_i$ and $\mathcal{U}_k$ separated by a vertical line, Barbay and Chen [3] described an algorithm that computes the common tangent between $\mathcal{U}_i$ and $\mathcal{U}_k$ in time within $O(\log a + \log b)$, where $\mathcal{U}_i[a]$ and $\mathcal{U}_k[b]$ are the points that lie in the tangent. At each

step this algorithm considers the points $\mathcal{U}_i[c]$ and $\mathcal{U}_k[d]$ and can certify at least in one upper hull sequence if the tangent is to the right or to the left of the point considered. A minor variant manages the case where the separating line is not vertical. Algorithm 2 executes several instances of this algorithm in parallel between $\mathcal{U}_j$ and all upper hull sequences to the right of p, always considering the same point in $\mathcal{U}_j$ (similarly to the Demaine et al.'s algorithm [10] to compute the intersection of sorted set). Once all parallel decisions about the point $\mathcal{U}_j[c]$ are made, the instances can be divided into two sets: (i) those whose tangents are to the left of $\mathcal{U}_j[c]$ and (ii) those whose tangents are to the right of $\mathcal{U}_j[c]$. Algorithm 2 stops the parallel computation of tangents in the set of maxima sequences (ii). The Step 5 continues until there is just one instance running and computes the tangent τ in this instance.

Analysis of the Quick Union Hull Algorithm. Similarly to the case of Merging Maxima, every algorithm for Merging Upper Hulls needs to certify that some blocks of the upper hull sequences can not participate in the Upper Hull of the union, and that some blocks of the upper hull sequences are in the Upper Hull of the union. In the following we formalize the notion of a *certificate* for Merging Upper Hulls problem.

Definition 6. *Given the points $\mathcal{U}_i[a]$ and $\mathcal{U}_j[b]$, let ℓ be the straight line that passes through $\mathcal{U}_i[a]$ and $\mathcal{U}_j[b]$ and let m_ℓ be the slope of ℓ. $\langle \mathcal{U}_i[a], \mathcal{U}_j[b] \supset \mathcal{U}_k[c..d..e] \rangle$ is an* Elementary Eliminator Argument *if all the points of the block $\mathcal{U}_k[c..e]$ are between the vertical lines through $\mathcal{U}_i[a]$ and $\mathcal{U}_j[b]$, $m(\mathcal{U}_k[d-1], \mathcal{U}_k[d]) \geq m_\ell \geq m(\mathcal{U}_k[d], \mathcal{U}_k[d+1])$, and the point $\mathcal{U}_k[d]$ lies below ℓ.*

If $\langle \mathcal{U}_i[a], \mathcal{U}_j[b] \supset \mathcal{U}_k[c..d..e] \rangle$ is an elementary eliminator argument then the points in the block $\mathcal{U}_k[c..e]$ can not participate in the Upper Hull of the union.

Lemma 8. *An elementary eliminator argument $\langle \mathcal{U}_i[a], \mathcal{U}_j[b] \supset \mathcal{U}_k[c..d..e] \rangle$ can be checked in constant time.*

Several blocks that are "eliminated" by the same pair of points can be combined into a single argument, a notion captured by the *block eliminator argument*.

Definition 7. *$\langle \mathcal{U}_i[a], \mathcal{U}_j[b] \supset \mathcal{U}_{k_1}[c_1..d_1..e_1], \ldots, \mathcal{U}_{k_t}[c_t..d_t..e_t] \rangle$ is a* Block Eliminator Argument *if $\langle \mathcal{U}_i[a], \mathcal{U}_j[b] \supset \mathcal{U}_{k_1}[c_1..d_1..e_1] \rangle, \ldots, \langle \mathcal{U}_i[a], \mathcal{U}_j[b] \supset \mathcal{U}_{k_t}[c_t..d_t..e_t] \rangle$ are elementary eliminator arguments.*

A block eliminator argument is checked by checking each elementary eliminator argument that form it.

Corollary 2. *A block eliminator argument $\langle \mathcal{U}_i[a], \mathcal{U}_j[b] \supset \mathcal{U}_{k_1}[c_1..d_1..e_1], \ldots, \mathcal{U}_{k_t}[c_t..d_t..e_t] \rangle$ can be checked in time within $O(t)$.*

As for Merging Maxima, any correct algorithm for Merging Upper Hulls must certify that some points are part of the output.

Definition 8. *$\langle \mathcal{U}_i[a] \dashv \mathcal{U}_{j_1}[b_1], \ldots, \mathcal{U}_{j_t}[b_t] \rangle$ is an* Elementary Convex Argument *if there exists a straight line ℓ that passes through $\mathcal{U}_i[a]$ of slope m_ℓ such that $m(\mathcal{U}_{j_1}[b_1 - 1], \mathcal{U}_{j_1}[b_1]) \geq m_\ell \geq m(\mathcal{U}_{j_1}[b_1], \mathcal{U}_{j_1}[b_1+1]), \ldots, m(\mathcal{U}_{j_t}[b_t-1], \mathcal{U}_{j_t}[b_t]) \geq m_\ell \geq m(\mathcal{U}_{j_t}[b_t], \mathcal{U}_{j_t}[b_t+1])$; $m(\mathcal{U}_i[a-1], \mathcal{U}_i[a]) \geq m_\ell \geq m(\mathcal{U}_i[a], \mathcal{U}_i[a+1])$; and the points $\mathcal{U}_{j_1}[b_1], \ldots, \mathcal{U}_{j_t}[b_t]$ lie below ℓ.*

If $\langle \mathcal{U}_i[a] \dashv \mathcal{U}_{j_1}[b_1], \ldots, \mathcal{U}_{j_t}[b_t] \rangle$ is an elementary convex argument, then the point $\mathcal{U}_i[a]$ is in the Upper Hull of the union of the upper hulls $\mathcal{U}_i, \mathcal{U}_{j_1}, \ldots, \mathcal{U}_{j_t}$.

Lemma 9. *An* elementary convex argument $\langle \mathcal{U}_i[a] \dashv \mathcal{U}_{j_1}[b_1], \ldots, \mathcal{U}_{j_t}[b_t] \rangle$ *can be checked in time within* $O(t)$.

There are blocks that can be "easily" certified that form part of the output.

Definition 9. *Given the points* $\mathcal{U}_i[a]$ *and* $\mathcal{U}_i[b]$*, let* ℓ *be the straight line that passes through* $\mathcal{U}_i[a]$ *and* $\mathcal{U}_i[b]$ *and let* m_ℓ *be the slope of* ℓ*.* $\langle \mathcal{U}_i[a..b] \dashv \mathcal{U}_{j_1}[c_1..d_1..e_1], \ldots, \mathcal{U}_{j_t}[c_t..d_t..e_t] \rangle$ *is a* Block Convex Argument *if* $\langle \mathcal{U}_i[a] \dashv \mathcal{U}_{j_1}[c_1], \ldots, \mathcal{U}_{j_t}[c_t] \rangle$ *and* $\langle \mathcal{U}_i[b] \dashv \mathcal{U}_{j_1}[e_1], \ldots, \mathcal{U}_{j_t}[e_t] \rangle$ *are elementary convex arguments;* $m(\mathcal{U}_{j_1}[d_1 - 1], \mathcal{U}_{j_1}[d_1]) \geq m_\ell \geq m(\mathcal{U}_{j_1}[d_1], \mathcal{U}_{j_1}[d_1 + 1]), \ldots, m(\mathcal{U}_{j_t}[d_t - 1], \mathcal{U}_{j_t}[d_t]) \geq m_\ell \geq m(\mathcal{U}_{j_t}[d_t], \mathcal{U}_{j_1}[d_t + 1])$*, and the points* $\mathcal{U}_{j_1}[d_1], \ldots, \mathcal{U}_{j_t}[d_t]$ *lie below* ℓ*.*

If $\langle \mathcal{U}_i[a..b] \dashv \mathcal{U}_{j_1}[c_1..d_1..e_1], \ldots, \mathcal{U}_{j_t}[c_t..d_t..e_t] \rangle$ is a block convex argument then the points in the block $\mathcal{U}_i[a..b]$ are in the Upper Hull of the union of the upper hulls $\mathcal{U}_i, \mathcal{U}_{j_1}, \ldots, \mathcal{U}_{j_t}$.

Lemma 10. *A* block convex argument $\langle \mathcal{U}_i[a..b] \dashv \mathcal{U}_{j_1}[c_1..d_1..e_1], \ldots, \mathcal{U}_{j_t}[c_t..d_t..e_t] \rangle$ *can be checked in time within* $O(t)$.

Similar to the Merging Maxima, the difficulty of finding and describing block eliminator and block convex arguments depend on the points they refer to in the upper hull sequences, a notion captured by "argument points":

Definition 10. *Given an argument* $\langle \mathcal{U}_i[a], \mathcal{U}_j[b] \supset \mathcal{U}_{k_1}[c_1..d_1..e_1], \ldots, \mathcal{U}_{k_t}[c_t..d_t..e_t] \rangle$ *or* $\langle \mathcal{U}_i[a..b] \dashv \mathcal{U}_{j_1}[c_1..d_1..e_1], \ldots, \mathcal{U}_{j_t}[c_t..d_t..e_t] \rangle$*, the* Argument Points *are the points* $\mathcal{U}_i[a]$ *and* $\mathcal{U}_i[b]$.

Those atomic arguments can be checked in time proportional to the number of blocks in them, and combine into a general definition of a certificate that any correct algorithm for Merging Upper Hulls in the algebraic decision tree computational model can be modified to output.

Definition 11. *Given a set of upper hull sequences and their* Upper Hull $\mathcal{U}$ *expressed as several blocks on the upper hull sequences. A* certificate *of* $\mathcal{U}$ *is a set of block eliminator and block convex arguments such that the* Upper Hull *of any instance satisfying those arguments is given by the description of* $\mathcal{U}$*. The length of a certificate is the number of distinct argument points in it.*

Similarly to the Merging Maxima, the key of the analysis is to separate the doubling search steps from the other steps of the algorithm.

Theorem 4. *Given* ρ *upper hull sequences. The time complexity of the* `Quick Union Hull` *algorithm is within* $O(\sum_{j=1}^{\beta} \log s_j + \sum_{i=1}^{\delta} \log \binom{\rho}{m_i})$ *when it computes the* Upper Hull *of the union of these upper hull sequences; where* β *is the number of blocks in the certificate* $\mathcal{C}$ *computed by the algorithm;* $s_1, \ldots, s_\beta$ *are the sizes of these blocks;* δ *is the length of* $\mathcal{C}$*; and* $m_1, \ldots, m_\delta$ *is a sequence where* m_i *is the number of upper hull sequences whose blocks form the* i*-th argument of* $\mathcal{C}$*.*

The optimality of this algorithm is a consequence of the fact that it checks each argument of any certificate using a constant number of argument points.

Lemma 11. *The algorithm* `Quick Union Hull` *computes a certificate which length is a constant factor of the length of the certificate of minimal length.*

In the following section we describe a synergistic results that combines the results of Sections 4.1 and 4.2.

4.3 Synergistic Computation of Upper Hulls

The `(Simple,Structure) Synergistic Hull` algorithm proceeds in two phases. It first decomposes the input into simple subchains of consecutive positions using a (new) linear time doubling search inspired partitioning algorithm, that searches for simple chains from left to right (see Algorithm 3 for a detail description of the algorithm). It computes their upper hull sequences, and then merges those using the `Quick Union Hull` algorithm described previously.

Algorithm 3 `Doubling Search Partition`

Input: A sequence of n planar points $p_1, \ldots, p_n$
Output: A sequence of simple polygonal chains
1: Initialize i to 1;
2: **for** $t = 1, 2, \ldots$ **do**
3: **if** $i + 2^t > n$ **or** the chain $p_i, \ldots, p_{i+2^t}$ is **not** simple **then**
4: Add the chain $p_i, \ldots, p_{i+2^{t-1}}$ to the output
5: Update $i \leftarrow i + 2^{t-1} + 1$
6: Reset $t \leftarrow 1$
7: **end if**
8: **end for**

The `Doubling Search Partition` algorithm partitions the polygonal chain into simple subchains which sizes has asymptotically minimum entropy among all the partitions into simple subchains. The following lemma formalizes this fact.

Lemma 12. *Given a sequence S of n planar points. The `Doubling Search Partition` algorithm computes in linear time a partition of S into k simple polygonal chain of consecutive points, of sizes $n_1, \ldots, n_k$ such that $n(1 + \mathcal{H}(n_1, \ldots, n_k)) \in O(n(1 + \alpha))$, where α is the minimal entropy $\min\{\mathcal{H}(n_1, \ldots, n_\kappa)$ of any partition of S into κ simple subchains of consecutive positions, of respective sizes $n_1, \ldots, n_\kappa$, and $\mathcal{H}(n_1, \ldots, n_\kappa) = \sum_{i=1}^{\kappa} \frac{n_i}{n} \log \frac{n}{n_i}\}$.*

The proof of this lemma is similar to the proof of Theorem 3 where the number of operations for each simple subchain of a partition into simple subchains is bounded separately. The following theorem summarizes the synergistic result in this section.

Theorem 5. *Let $\mathcal{S}$ be a set of points in the plane such that $\mathcal{S}$ can be partitioned into κ simple subchains. Let h be the number of points in the* Upper Hull *of $\mathcal{S}$. There exists an algorithm which time complexity is within $O(n + \sum_{i=1}^{\delta} \log \binom{\kappa}{m_i}) \subseteq O(n \log(\min(\kappa, h)))$ when it computes the* Upper Hull *of $\mathcal{S}$, where δ is the length of the certificate $\mathcal{C}$ computed by the union algorithm; and $m_1, \ldots, m_\delta$ is a sequence where m_i is the number of upper hull sequences of the simple subchains whose blocks form the i-th argument of $\mathcal{C}$. This number of comparisons is optimal in the worst case over instances $\mathcal{S}$ formed by κ simple subchains which* Upper Hulls *have certificates $\mathcal{C}$ of length δ such that $m_1, \ldots, m_\delta$ is a sequence where m_i is the number of upper hull of simple subchains whose blocks form the i-th argument of $\mathcal{C}$.*

Even though the algorithms are more complex, except for some details, the proofs of Theorems 4 and 5 and Lemma 11 of Section 4 are very similar to those described in the previous section. We describe the intuition for the lower bound below: as for the computation of Maxima Sets, it is a simple adversary argument, based on the definition of a family of "hard" instances for each possible value of the parameters of the analysis, building over each other, but the combination of elementary instances requires a little bit of extra care.

First, we verify the lower bound for "easy" instances, of finite difficulty: general instances formed by a single ($\kappa = 1$) simple sequence obviously require $\Omega(n)$ comparisons (no correct algorithm can afford to ignore a single point of the input), while general instances dominated by a single edge ($h = 1$) also require $\Omega(n)$ comparisons. Each of these lower bound yields a distribution of instances, either decomposed into $\kappa = 1$ simple chains or of output size $h = 1$, such that any deterministic algorithm performs $\Omega(n)$ comparisons on average on a uniform distribution of those instances.

Such distributions of "elementary" instances can be duplicated so that to produce various distributions of elementary instances; and can be combined so that to define a distribution of harder instances.

Lemma 13. *Given the positive integers $n, \kappa, \beta, s_1, \ldots, s_\beta, \delta, m_1, \ldots, m_\delta$, there is a family of instances which can each be partitioned into κ simple subchains such that,*

- *β and $s_1, \ldots, s_\beta$ are the number and sizes of the blocks in the certificate $\mathcal{C}$ computed by the union algorithm, respectively;*
- *δ is the length of $\mathcal{C}$;*
- *$m_1, \ldots, m_\delta$ is a sequence where m_i is the number of upper hulls of the simple subchains whose blocks form the i-th argument of $\mathcal{C}$; and*
- *on average on a uniform distribution of these instances, any algorithm computing the* Upper Hull *of S in the comparison model performs within $\Omega(n + \sum_{i=1}^{\delta} \log \binom{\kappa}{m_i})$ comparisons.*

Finally, any such distribution with a computational lower bound on average yields a computational lower bound for the worst case instance complexity of any randomized algorithm, on average on its randomness; and as a particular case a lower bound on the worst case complexity of any deterministic algorithm:

Corollary 3. *Given the positive integers $\kappa, \beta, s_1, \ldots, s_\beta, \delta, m_1, \ldots, m_\delta$, and an algorithm A computing the* Upper Hull *of a sequence of n planar points in the algebraic decision tree computational model (whether deterministic or randomized), there is an instance I such that A performs a number of comparisons within $\Omega(n + \sum_{i=1}^{\delta} \log \binom{\kappa}{m_i})$ when it computes the* Upper Hull *of I.*

Proof. A direct application of Yao's minimax principle [20, 23, 25]. □

This concludes the description of our synergistic results. In the next section, we discuss the issues left open for improvement.

5 Discussion

Considering the computation of the Maxima Set and of the Convex Hull, we have built upon previous results taking advantage either of some notions of input order or of some notions of input structure, to describe solutions which take advantage of both in a synergistic way. There are many ways in which those results can be improved further: we list only a selection here.

In the same line of thought, Ahn and Okamoto [2] described some other notion of input order than the one we considered here, which can potentially yield another synergistic solution in combination with a given notion of input structure. This is true for any of the many notions of input order which could be adapted from Sorting [18].

Whereas being adaptive to as many measures of difficulty as possible at once is a worthy goal in theory, it usually comes at a price of an increase in the constant factor of the running time of the algorithm: it will become important to measure, for the various practical applications of each

problem, which measures of difficulty take low value in practice. It will be necessary to do some more theoretical work to identify what to look for in the practical applications, but then it will be important to measure the practical difficulties of the instances.

Acknowledgments: The authors would like to thank Javiel Rojas for helping with the bibliography on the computation of the Maxima Set of a set of points.

References

1. Afshani, P., Barbay, J., Chan, T.M.: Instance-optimal geometric algorithms. In: Proceedings of the Annual IEEE Symposium on Foundations of Computer Science (FOCS). pp. 129–138. IEEE Computer Society (2009)
2. Ahn, H.K., Okamoto, Y.: Adaptive algorithms for planar convex hull problems. IEICE Transactions 94-D(2), 182–189 (2011)
3. Barbay, J., Chen, E.Y.: Convex hull of the union of convex objects in the plane: an adaptive analysis. In: Proceedings of the Annual Canadian Conference on Computational Geometry (CCCG) (2008)
4. Barbay, J., Navarro, G.: On compressing permutations and adaptive sorting. Theoretical Computer Science (TCS) 513, 109–123 (2013)
5. Barbay, J., Ochoa, C., Rao, S.S.: Synergistic Sorting and Deferred Data Structures on MultiSets. ArXiv e-prints (Aug 2016)
6. Bentley, J.L., Yao, A.C.C.: An almost optimal algorithm for unbounded searching. Information Processing Letters (IPL) 5(3), 82–87 (1976)
7. Chan, T.M., Snoeyink, J., Yap, C.K.: Primal dividing and dual pruning: Output-sensitive construction of four-dimensional polytopes and three-dimensional voronoi diagrams. Discrete & Computational Geometry (DCG) 18(4), 433–454 (1997)
8. Chazelle, B.: Triangulating a simple polygon in linear time. Discrete & Computational Geometry (DCG) 6(5), 485–524 (Aug 1991), http://dx.doi.org/10.1007/BF02574703
9. Demaine, E.D., Lopez-Ortiz, A.: A linear lower bound on index size for text retrieval. In: Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete algorithms (SODA). pp. 289–294 (2001)
10. Demaine, E.D., López-Ortiz, A., Munro, J.I.: Adaptive set intersections, unions, and differences. In: Proceedings of the 11^{th} ACM-SIAM Symposium on Discrete Algorithms (SODA). pp. 743–752 (2000)
11. Estivill-Castro, V., Wood, D.: A survey of adaptive sorting algorithms. ACM Computing Surveys (ACMCS) 24(4), 441–476 (1992)
12. Fjällström, P.O., Katajainen, J., Levcopoulos, C., Petersson, C.: A sublogarithmic convex hull algorithm. BIT Numerical Mathematics 30(3), 378–384 (1990)
13. Kirkpatrick, D.G., Seidel, R.: Output-size sensitive algorithms for finding maximal vectors. In: Proceedings of the Annual Symposium on Computational Geometry (SoCG). pp. 89–96. ACM, New York, NY, USA (1985)
14. Kirkpatrick, D.G., Seidel, R.: The ultimate planar convex hull algorithm? SIAM Journal on Computing (SICOMP) 15(1), 287–299 (1986)
15. Kung, H.T., Luccio, F., Preparata, F.P.: On finding the maxima of a set of vectors. Journal of the ACM (JACM) 22(4), 469–476 (1975)
16. Levcopoulos, C., Lingas, A., Mitchell, J.S.B.: Adaptive algorithms for constructing convex hulls and triangulations of polygonal chains. In: Proceedings of the Scandinavian Workshop on Algorithm Theory (SWAT). pp. 80–89. Springer-Verlag, London, UK (2002)
17. Melkman, A.A.: On-line construction of the convex hull of a simple polyline. Information Processing Letters (IPL) 25(1), 11–12 (Apr 1987), http://dx.doi.org/10.1016/0020-0190(87)90086-X
18. Moffat, A., Petersson, O.: An overview of adaptive sorting. Australian Computer Journal (ACJ) 24(2), 70–77 (1992)
19. Munro, J.I., Spira, P.M.: Sorting and searching in multisets. SIAM Journal on Computing (SICOMP) 5(1), 1–8 (1976)
20. Neumann, J.V., Morgenstern, O.: Theory of games and economic behavior. 1st edition, Princeton University Press (1944)

21. Preparata, F.P., Hong, S.J.: Convex hulls of finite sets of points in two and three dimensions. Communication of the ACM (CACM) 20, 87–93 (1977)
22. Preparata, F.P., Shamos, M.I.: Computational Geometry: An Introduction. Springer-Verlag (1985)
23. Sion, M.: On general minimax theorems. Pacific Journal of Mathematics (PJM) 1, 171–176 (1958)
24. Takaoka, T.: Partial solution and entropy. In: Královič, R., Niwiński, D. (eds.) Mathematical Foundations of Computer Science (MFCS) 2009: 34th International Symposium, Novy Smokovec, High Tatras, Slovakia, August 24-28, 2009. Proceedings. pp. 700–711. Springer Berlin Heidelberg, Berlin, Heidelberg (2009)
25. Yao, A.C.C.: Probabilistic computations: Toward a unified measure of complexity. In: Proceedings of the Annual IEEE Symposium on Foundations of Computer Science (FOCS). pp. 222–227 (1977)